

ParlayMarket: Automated Market Making for Parlay-Style Joint Contracts

Ranvir Rana  

Kaleidoscope Blockchain Inc., Raleigh, NC, USA

Viraj Nadkarni   

Princeton University, NJ, USA

Niusha Moshrefi

Princeton University, NJ, USA

Pramod Viswanath

Princeton University, NJ, USA

Abstract

Prediction markets are powerful mechanisms for information aggregation, but existing designs are optimized for single-event contracts. Traders frequently express beliefs about joint outcomes – sports parlays, conditional forecasts, multi-scenario financial bets. Current platforms either prohibit such trades or rely on ad hoc mechanisms that ignore correlation structure, resulting in inefficient prices and fragmented liquidity. We introduce *ParlayMarket*, an automated market-maker for parlay-style joint contracts. The mechanism maintains a shared pairwise exponential-family belief state, so all base and parlay prices are marginals of one coherent distribution. This compresses the 2^M outcome space into $O(M^2)$ sufficient statistics and allows one liquidity pool to support an exponentially large family. Our main result characterizes the resulting learning and loss dynamics. Under repeated trading, prices converge to the best pairwise approximation of the true joint distribution. The induced expected market-maker loss grows at most quadratically in the number of base events, rather than exponentially in the number of listed parlays; moreover, this quadratic dependence is worst-case optimal for dense pairwise dependence, since there are $O(M^2)$ independent correlation directions to learn. Parlay trades are essential to this guarantee: they provide direct constraints on joint outcomes and reduce steady-state error relative to learning from marginal trades alone. Experiments on synthetic correlated markets and historical Kalshi combo data confirm the predicted scaling and show that the mechanism remains effective in realistic market-making settings. Our results demonstrate that combinatorial expressiveness does not require combinatorial capital.

2012 ACM Subject Classification Theory of computation → Market equilibria; Theory of computation → Online learning theory; Applied computing → Economics

Keywords and phrases Prediction markets, automated market makers, parlays, market scoring rules, online learning

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.15

Related Version *Full Version*: <https://arxiv.org/abs/2603.22596>

1 Introduction

Two uses of knowledge in a society. Knowledge in an economy can be used in two ways – *factually* and *counterfactually*. Factual knowledge concerns statements such as “the price of oil will be greater than \$100 on March 29” or “the election will be won by party P .” Counterfactual knowledge, by contrast, concerns statements such as “if the price of oil is greater than \$100 on March 29, then the price of a banana will be greater than \$2 on April 30” or “if the election is won by party P , then the probability of a recession will increase.” Prediction markets have proven to be a powerful mechanism for aggregating



© Ranvir Rana, Viraj Nadkarni, Niusha Moshrefi, and Pramod Viswanath;
licensed under Creative Commons License CC-BY 4.0

8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyrpoulou; Article No. 15; pp. 15:1–15:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

and eliciting factual knowledge. The prices quoted by a prediction market can serve as a public good, allowing uninformed members of the economy to peer into the future through a market-implied best guess and plan accordingly.

The liquidity problem. Despite this promise, prediction-market liquidity is highly uneven. Existing markets tend to have substantial liquidity in sensational or gambling-oriented domains such as sports, while economically useful markets, including public policy, local politics, and macroeconomic indicators, are often woefully illiquid [9]. Automated market makers (AMMs) are particularly well suited to such illiquid markets because they can provide continuous liquidity without requiring a counterparty [11, 12]. However, conventional AMMs incur losses on their locked liquidity in exchange for eliciting information. In this sense, subsidies are required to induce traders to reveal factual knowledge.

The joint events problem. A further challenge arises when the information of interest is not merely factual, but counterfactual. Prediction markets aggregate dispersed information into prices, but practical automated market makers are primarily designed for single-event contracts. In particular, bounded-loss pricing for single markets is typically achieved through Hanson’s logarithmic market scoring rule (LMSR) [11, 12]. In many applications, however, traders reason about joint outcomes: sports parlays, conditional forecasts across related events, or multi-scenario financial bets. The probability of such a contract depends not only on the marginal probabilities of its constituent events, but also on the dependence structure among them.

Computational complexity. Supporting such joint contracts introduces a fundamental challenge. With M binary events, the joint outcome space contains 2^M possible realizations. A naive combinatorial market maker must maintain prices and risk exposure across all of these outcomes, leading to worst-case loss and computational requirements that scale exponentially in M . This renders direct approaches intractable beyond very small systems. The challenge is well known in combinatorial prediction markets: allowing direct trading over joint outcome spaces provides expressive contracts, but typically incurs exponential state and inference costs [12, 8]. Existing practical systems therefore either avoid parlay-style contracts or handle them through ad hoc mechanisms such as request-for-quote workflows, which fragment liquidity and need not maintain a coherent joint distribution.

Our contributions. This paper shows that *combinatorial expressiveness does not require combinatorial capital*. We introduce *ParlayMarket*, an automated market maker that maintains a shared probabilistic belief over all outcomes and derives every base and parlay price as a marginal of this belief. The key design choice is to restrict the belief state to a pairwise exponential family. Rather than tracking one state variable per joint outcome, the mechanism maintains only

$$M + \binom{M}{2} = O(M^2)$$

sufficient statistics: one marginal parameter per event and one interaction parameter per pair. This is the smallest generic exponential-family representation capable of encoding arbitrary pairwise dependence: independence cannot represent correlations, while an unrestricted joint model is exponential. Statistically, this model is also the maximum-entropy distribution matching the learned first and second moments, and therefore does not introduce higher-order structure unsupported by the observed trade flow [25].




Figure 1 ParlayMarket enables 2^M markets with $O(M^2)$ capital. Theoretical curves show the asymptotic loss scaling implied by the analysis in Section 5. Empirical curves are obtained from controlled simulation of correlated binary markets under the Gaussian score model described later in Sections 3 and 6, and confirm the same polynomial-vs.-exponential separation in practice.

Our approach is related in spirit to inference-based views of prediction markets, where prices are treated as marginals of an underlying probabilistic model [10, 8, 17]. The key difference is that ParlayMarket learns the dependence structure online from endogenous order flow rather than imposing a fixed graphical structure *ex ante*. Trades in both base markets and parlay markets are interpreted as observations about the underlying joint distribution. After each trade, the mechanism updates its shared parameters, allowing information from one contract to propagate immediately to related contracts. This couples liquidity and learning across the entire contract family rather than maintaining separate books for each parlay.

We show that this design satisfies three central objectives. First, it pools liquidity across markets by pricing all contracts from a shared state, avoiding fragmentation. Second, it achieves bounded loss through an LMSR-style cost function defined over this shared representation, with loss that grows at most quadratically in M , rather than exponentially in the number of joint outcomes. Moreover, this quadratic dependence is worst-case optimal for dense pairwise dependence. Third, it aggregates information efficiently: as trades accumulate, the model converges to the best pairwise approximation of the true joint distribution under the observed data. This does not claim that higher-order dependence is absent; rather, it identifies the pairwise family as the practical frontier at which global coherence, online learnability, and non-combinatorial loss can all be achieved simultaneously.

Conceptually, ParlayMarket transforms the market into an online inference system. Empirically, controlled simulations confirm the predicted polynomial-versus-exponential loss separation, as shown in Figure 1, and a replay on historical Kalshi combo data shows that the mechanism remains effective as a practical correlation-aware quoting engine. Together, these results show that parlay-style prediction markets can be supported by an automated market maker without full joint enumeration or separate capitalization for every combination.

Paper Overview. The remainder of the paper is organized as follows. In Section 2, we review background on prediction markets, combinatorial and parlay markets, and probabilistic models of dependence, and position our work relative to prior approaches. In Section 3, we formalize the problem and define the design objectives for parlay-capable market makers. Section 4 analyzes existing approaches and shows that they fail to simultaneously satisfy these objectives due to fragmented liquidity and combinatorial loss. Section 5 introduces ParlayMarket and

describes its pricing and update mechanisms. We then present the theoretical results on convergence, loss bounds, and information aggregation. Finally, Section 6 evaluates the mechanism empirically, and Section 7 discusses implications and future directions. Additional proofs, experimental results, and implementation details are provided in the full version of the paper.

2 Background and Related Work

Cost-function market makers. Prediction markets are commonly implemented using cost-function market makers, most notably Hanson’s logarithmic market scoring rule (LMSR) [12, 13]. LMSR provides continuous liquidity, proper-scoring-rule prices, and bounded worst-case loss. Chen and Pennock [7] interpret LMSR as a utility-based market maker, while Abernethy et al. [1] connect market scoring rules to convex optimization, online learning, and Bregman divergence minimization. Related work on exponential-family markets [2] shows how bounded-loss market makers can be understood as maintaining beliefs in a parametric family. CFMMs provide a parallel AMM tradition in decentralized finance [3, 19], but these mechanisms do not by themselves enforce probabilistic consistency across related event contracts.

Combinatorial and Parlay Markets. Combinatorial prediction markets allow trading over joint outcome spaces [12], but exact implementations require reasoning over 2^M outcomes and become intractable for large M . Prior work therefore restricts the security space or market geometry to obtain tractable loss bounds [15], or studies settings where simple trader interfaces fail to recover coherent joint beliefs [20]. In practice, platforms such as Polymarket and Kalshi [21, 16] rarely implement full combinatorial AMMs; parlay-like products are typically handled through RFQ workflows or separate markets, which fragment liquidity and need not maintain a single coherent joint distribution.

Graphical models and inference-based markets. Graphical models offer a natural way to represent dependence compactly. Pairwise Markov random fields, including Ising models, represent binary joint distributions using only first-order and pairwise interaction terms and correspond to maximum-entropy distributions matching specified marginals and pairwise moments [25]. Learning such models has been studied extensively [22, 5, 24], and online variants connect naturally to stochastic optimization and online convex programming [26, 14]. Prediction markets have also been interpreted as probabilistic inference systems [10, 6]; Bayesian-network and graphical-model market makers make this connection explicit [8, 17]. Unlike those mechanisms, *ParlayMarket* learns the dependence structure endogenously from singleton and parlay order flow while using one shared state to quote all contracts coherently.

3 Model and Design Objectives

We formalize automated market making for correlated prediction markets as a joint mechanism design and statistical inference problem. The market maker interacts with a stochastic stream of trades that reveal partial information about an underlying joint distribution, and must simultaneously provide liquidity, control risk, and learn the dependence structure.

3.1 Market Setting

We consider a collection of M binary events $X_1, \dots, X_M \in \{0, 1\}$, and denote a joint outcome by $x \in \{0, 1\}^M$. The uncertainty over outcomes is governed by an unknown ground-truth distribution P^* over $\{0, 1\}^M$, which captures both marginal probabilities and dependencies among events.

Contracts. The market supports trading in a family of payoff functions indexed by subsets $S \subseteq \{1, \dots, M\}$. Each contract corresponds to a Boolean payoff

$$f_S(x) = \prod_{i \in S} x_i \quad (1)$$

which pays 1 if all events in S occur and 0 otherwise. In particular, singleton sets $S = \{i\}$ recover base contracts, while larger subsets correspond to parlay contracts. Let $\mathcal{F}$ denote the set of all supported contracts.

Pricing and coherence. At any time, the market maker maintains a parametric distribution P_φ over $\{0, 1\}^M$, which induces prices for all contracts via expectations:

$$p_S(\varphi) = \mathbb{E}_{x \sim P_\varphi}[f_S(x)] \quad (2)$$

This ensures probabilistic coherence, in the sense that all prices are consistent with a single joint distribution. In particular, marginal and joint prices satisfy the usual consistency constraints implied by P_φ .

3.2 Trader/Information Model

We model the market as an online information aggregation process in which traders arrive sequentially and trade based on private signals about the underlying outcome distribution.

Underlying process. The underlying uncertainty in each base market is driven by a correlated Gaussian score process. Specifically, there are M base markets where market i resolves YES if an underlying score $S_i(T)$ exceeds threshold K_i at expiry T . The scores follow correlated arithmetic Brownian motion

$$dS_i(t) = \sigma_i dW_i(t), \quad \text{Cov}(dW_i, dW_j) = \rho_{ij} dt, \quad (3)$$

so that $S_T | S_t \sim \mathcal{N}(S_t, \Sigma)$ with $\Sigma_{ij} = \rho_{ij} \sigma_i \sigma_j (T - t)$. The binary outcome of market i is $X_i = \mathbf{1}[S_i(T) > K_i]$.

The Gaussian score model is a natural choice for several reasons. First, it provides a tractable closed-form expression for the conditional probability of each outcome given the current score, enabling informed traders to compute exact signals. Second, the pairwise correlation structure ρ_{ij} maps cleanly onto the interaction parameters W_{ij} of the Ising approximation used by ParlayMarket, making the theoretical analysis transparent. Third, it is a standard model for continuous-time prediction markets and sports scoring processes, and admits straightforward emulation from real-world data by calibrating $(\sigma_i, \rho_{ij}, K_i)$ to observed outcome frequencies and co-occurrences [23]. In particular, it captures how the volatility in the likelihood of an event increases as we move closer to when the outcome is going to be revealed.

Trader types. At each time step, a trader arrives and interacts with a single contract. *Informed traders* observe the true underlying score process S_t and target the exact conditional probability. For base market i they submit $\hat{p}_i = \Phi\left(\frac{S_t^i - K_i}{\sigma_i \sqrt{T-t}}\right)$, and for parlay $\mathcal{S}$ they submit the exact multivariate normal orthant probability $P(\bigcap_{i \in \mathcal{S}} \{S_i(T) > K_i\} \mid S_t)$. *Noise traders* observe a noisy signal $\tilde{S}_t^i = S_t^i + \varepsilon^i$ with $\varepsilon \sim \mathcal{N}(0, \text{diag}(\sigma_i^2(T-t)))$ independent of the path, and target the resulting noisy probability estimate. The noise-trader fraction is $\alpha \in [0, 1]$.

Information Structure. Each trader has access only to a local signal about a bounded number of events – there is no global information. Informed traders observe the score process $S_t \in \mathbb{R}^M$ for all M base markets, but this is a signal about the current conditional distribution only; no trader has access to the full future path or to the joint distribution P^* directly. PM itself must learn P^* online from the sequence of incoming trades.

3.3 Market Interface

The market maker exposes a pricing and trading interface over the contract set $\mathcal{F}$. Trades are executed via a cost function $C_S(\cdot; \theta)$, which determines the payment required to adjust the position in contract S . Following each trade, the market maker updates its internal parameters θ , which in turn updates prices across all contracts.

Interaction model. Traders arrive sequentially. At each round t , one market m_t is selected uniformly at random from the $M + P$ markets, and a single trader arrives. With probability $1 - \alpha$ the arriving trader is informed; with probability α they are a noise trader. Competition is fair at each step: every arriving trader observes the current PM price and decides whether and how much to trade before the next trader arrives.

3.4 Design Objectives

We formalize the requirements for a market maker supporting correlated contracts.

Property 1: Liquidity Consistency. The market maker should provide comparable liquidity across all contracts without requiring independent capital allocation.

Let $I_m(x; \phi)$ denote the price impact of a trade x in market m . Then there should exist constants $0 < c_1 \leq c_2 < \infty$ such that for all sufficiently small admissible trade sizes x , such that there is non-zero price impact across all markets,

$$c_1 \leq \frac{I_m(x; \phi)}{I_{m'}(x; \phi)} \leq c_2 \quad \forall m, m' \in \mathcal{M} \quad (4)$$

That is, price impact should remain comparable across all markets, including both base and parlay contracts, so that parlays are not disproportionately illiquid.

Property 2: Controlled loss. The market maker should control expected aggregate loss and tail risk under the induced joint process of trades and outcomes.

Let L_T denote cumulative realized AMM loss over all executed trades up to horizon T . A satisfactory mechanism should control both expected aggregate loss and tail risk under the induced joint process of trades and outcomes.

First, expected loss should scale subexponentially in the number of base markets, ideally linearly:

$$\mathbb{E}[L_T] = o(2^M), \quad \text{ideally } \mathbb{E}[L_T] = O(M) \quad (5)$$

Second, tail loss should also remain controlled. For a confidence level α such as 0.95, define

$$\text{VaR}_\alpha(L_T) := \inf\{\ell : \Pr(L_T \leq \ell) \geq \alpha\} \quad (6)$$

Then the mechanism should satisfy

$$\text{VaR}_{0.95}(L_T) = o(2^M), \quad \text{ideally } \text{VaR}_{0.95}(L_T) = O(M) \quad (7)$$

Optionally, the same requirement can be imposed on $\text{CVaR}_{0.95}(L_T)$.

Property 3: Information aggregation. The market maker should aggregate information from trades so that the implied distribution converges to the true distribution P^* .

Let P_{ϕ_T} denote the distribution represented by the AMM after T trades. A satisfactory mechanism should use singleton and parlay order flow to recover the latent joint law with small statistical error.

In particular, after a number of trades linear in the effective model dimension, the represented distribution should be close to the true distribution in KL divergence:

$$D_{\text{KL}}(P^* \parallel P_{\phi_T}) \leq \varepsilon \quad \text{for } T = O(n), \quad (8)$$

where n denotes the relevant model dimension or number of informative parameters. Equivalently, the KL error should decrease with the number of informative trades:

$$\mathbb{E}[D_{\text{KL}}(P^* \parallel P_{\phi_T})] \rightarrow 0 \quad \text{as } T \text{ grows.} \quad (9)$$


Implication for market design. Taken together, Properties 1–3 indicate that neither pure contract-by-contract execution nor pure contract-by-contract pricing is sufficient. Treating each exposed parlay as an isolated market tends to fragment liquidity and scale risk with the number of listed contracts, while quoting combinations purely as functions of current base-market prices can improve accessibility but does not, by itself, recover the latent joint law. The design problem is therefore to support executable parlay trading while pooling risk and information strongly enough to preserve liquidity, control aggregate loss, and use order flow to learn dependence. The next section examines two natural baseline approaches along these lines and shows where each falls short.

4 Baseline Designs for Parlay-style Markets

The properties defined in Section 3 impose stringent requirements on any viable market design. We examine two natural baseline approaches – request-for-quote systems and independence-based automated market makers – and show that each fails to satisfy these requirements for fundamental, structural reasons. Additional empirical details, robustness checks, and full raw results are provided in the full version.

4.1 Request-for-Quote Protocol

In many existing platforms, parlay trades are handled via a request-for-quote (RFQ) mechanism: a trader specifies a bundle of outcomes, and a market maker returns a price based on proprietary models. This architecture is flexible, but it externalizes both pricing and risk management to separate dealers. As a result, quotes are not necessarily derived from a common joint distribution, and liquidity is fragmented across requested combinations.



■ **Figure 2** The RFQ trilemma.

■ **Table 1** Average spreads in base and combo markets on Kalshi, restricting to markets with open interest between 1,000 and 2,000. Combo markets exhibit substantially wider spreads than base markets.

Market type	Number of markets	Average spread (cents)
Base markets	999	5.12
Combo markets	549	11.89

Shortcomings of Using RFQ Protocol for Parlays. This creates two structural failures. First, liquidity deteriorates for combinations: in our Kalshi combo sample, average spreads are substantially larger in combo markets than in base markets (Table 1). Second, RFQ quotes need not be probabilistically coherent. Since quotes are not marginals of a shared joint law, they can violate basic constraints such as the Fréchet bound $P(A \cap B) \leq \min\{P(A), P(B)\}$; detailed violations are reported in Appendix B.1 of the full version. Aggregating quotes from multiple market makers yields only marginal improvements in spread, suggesting that fragmentation persists because competition among quote providers remains limited. A detailed breakdown by leg count and market is provided in Appendix B.2 of the full version.

The RFQ trilemma. The RFQ architecture faces a fundamental trilemma among the three properties introduced in Section 3. Opening maker participation can improve quote availability and price discovery, but weakens loss control because risk is dispersed across less tightly capitalized makers. Conservative quoting protects makers from adverse selection, but the resulting quotes become poor probability estimates and therefore carry little information. Restricting the set of quoted combinations makes risk easier to manage, but sacrifices broad parlay liquidity. Thus, over a broad parlay family, RFQ systems can trade off among liquidity consistency, coherent information aggregation, and bounded loss, but cannot achieve all three simultaneously. We call this structural constraint the *RFQ trilemma* (Figure 2).

4.2 Independence-Based Parlay AMM

A second natural baseline is to keep independent AMMs for the base markets and quote a parlay as the product of its marginal prices:

$$q_{\text{ind}}(E_1 \cap \cdots \cap E_k) = \prod_{i=1}^k \hat{p}_i. \quad (10)$$

This gives immediate executable quotes, but it hard-codes independence and therefore fails whenever outcomes are correlated.

The failure is structural. Consider two positively correlated events with common marginal probability p . The independence quote for the joint contract is p^2 , while the true joint probability can be close to p . An informed trader can repeatedly buy the parlay at p^2 against expected payoff p . Because parlay trades do not update the base-market prices or any shared correlation state, the mispricing does not disappear through trading. The AMM therefore fails both bounded loss and information aggregation: it cannot learn the dependence structure no matter how much parlay order flow it observes.

4.3 Towards a Joint Distribution Market Maker

These observations suggest that any viable design must satisfy three structural requirements: (i) maintain a shared representation of the joint distribution, (ii) update this representation from both base and composite trades, and (iii) derive all prices coherently from this shared state. The next section introduces ParlayMarket, an AMM that satisfies these requirements by maintaining a tractable approximation to the joint distribution and updating it via online learning from trade flow.

5 ParlayMarket

The limitations of baseline designs in Section 4 arise from the absence of a shared global state and the inability to propagate information across markets. We now construct a market maker that resolves these limitations by maintaining and updating a coherent joint distribution over all outcomes.

At a high level, ParlayMarket interprets each trade – whether in a base market or a parlay – as a noisy observation about the underlying joint distribution P^* . The mechanism maintains a parametric approximation P_ϕ to this distribution, derives all prices as marginals of P_ϕ , and updates ϕ after every trade. This transforms the market into an online inference system, where prices reflect a globally consistent and continuously learned probabilistic model.

We show that this design satisfies all three properties from Section 3, and that the complete-parlay variant enjoys an additional structural bonus: per-market loss halves with each new base asset added.

5.1 Architecture

The central challenge for a multi-market maker is keeping prices across all $N_M = 2^M - 1$ markets mutually consistent as information arrives through trade flow. A trade on any single market reveals something about the joint distribution, and a rational AMM should propagate this information to related markets immediately – otherwise prices become stale and the AMM is exposed to arbitrage. The ParlayMarket solves this through three interlocking mechanisms: *shadow trades* that broadcast information across markets, an *Ising model* belief state that compactly represents the joint distribution, and a *composite likelihood SGD* update that continuously calibrates the model from observed trade flow.

Shadow trades and price propagation. Prices in the ParlayMarket are posted marginals of a shared internal belief distribution P_ϕ over all M outcomes. When a trader bets on market S , the observed trade direction reveals information not just about event S but about all correlated events. Since we assume all incoming traders are informed (Section 3), the AMM should treat each trade as a signal about the entire joint distribution and update prices everywhere, not only in market S .

To implement this, the AMM records a set of internal *shadow trades* on related markets after each real trade arrives. Shadow trades are not real bets – they are bookkeeping entries that encode the information the AMM infers about other markets from the observed trade. Critically, they ensure that no market’s price becomes stale simply because it receives little direct trading volume; even an esoteric parlay market that rarely attracts traders will see its posted price updated whenever a correlated market is active.

Ising model as the joint belief state. To propagate information through shadow trades, the AMM needs a tractable model of the joint distribution that can be updated incrementally. Parlay trades are especially informative here: a bet on a k -way parlay directly reveals a joint probability $P^*(\bigcap_{i \in S} \{X_i = 1\})$, which constrains not only the individual event probabilities but also their pairwise and higher-order correlations.

The Ising model is the natural parametric choice. Given a target set of marginal probabilities $\{p_i^*\}$ and pairwise correlations $\{p_{ij}^*\}$, the Ising model is the *maximum-entropy* distribution consistent with those first and second-order moments:

$$P_{\boldsymbol{\varphi}}(\mathbf{x}) \propto \exp \left(\sum_{i=1}^M \theta_i x_i + \sum_{i < j} W_{ij} x_i x_j \right), \quad (11)$$

where $\boldsymbol{\varphi} = (\boldsymbol{\theta}, \mathbf{W})$ collects bias parameters $\theta_i \in \mathbb{R}$ (controlling individual event probabilities) and interaction weights $W_{ij} \in \mathbb{R}$ (controlling pairwise correlations). Choosing the maximum-entropy distribution is equivalent to making the weakest possible assumptions beyond the moment constraints: no spurious higher-order structure is imposed that is not supported by the data. This also corresponds to a moment-matching approach to Bayesian updating – rather than tracking the full 2^M -dimensional joint distribution, the AMM only needs to maintain the $\binom{M}{2} + M$ sufficient statistics ($\mathbb{E}[X_i], \mathbb{E}[X_i X_j]$) of the Ising exponential family.

Posted prices are the Ising marginals computed via belief propagation: $p_i^{\boldsymbol{\varphi}} = P_{\boldsymbol{\varphi}}(X_i = 1)$ for base markets and $p_S^{\boldsymbol{\varphi}} = P_{\boldsymbol{\varphi}}(\bigcap_{i \in S} \{X_i = 1\})$ for a parlay $S \subseteq [M]$. All N_M prices are generated from the single shared parameter vector $\boldsymbol{\varphi}$, so they are guaranteed to be globally consistent with some joint distribution at every point in time.

Gradient updates from trade errors. Each trade on market m reveals a discrepancy between the AMM’s current posted price $p_m^{\boldsymbol{\varphi}}$ and the true probability $p_m^* = P^*(\mathcal{E}_m)$ implied by the trader’s signal. The AMM treats this discrepancy as a cross-entropy error and takes a stochastic gradient step to reduce it. Formally, the AMM implicitly minimises the *composite pseudo-likelihood* [4, 18]:

$$\mathcal{L}(\boldsymbol{\varphi}) = \sum_{i=1}^M \lambda_i \text{CE}(p_i^*, p_i^{\boldsymbol{\varphi}}) + \sum_{i < j} \lambda_{ij} \text{CE}(p_{ij}^*, p_{ij}^{\boldsymbol{\varphi}}), \quad (12)$$

where $\lambda_i, \lambda_{ij} \geq 0$ are normalised trade arrival rates and $\text{CE}(p, q) = -p \log q - (1-p) \log(1-q)$. Upon each trade on market m , the AMM applies the SGD update

$$\boldsymbol{\varphi}_{t+1} = \boldsymbol{\varphi}_t - \eta \nabla_{\boldsymbol{\varphi}} \text{CE}(p_m^*, p_m^{\boldsymbol{\varphi}_t}), \quad (13)$$

with step size $\eta > 0$. Because $\boldsymbol{\varphi}$ is shared across all markets, a gradient step triggered by a trade on market m immediately shifts the posted prices of every other market through the Ising model’s interaction terms. Shadow trades are then computed using the updated $\boldsymbol{\varphi}_{t+1}$, so the propagated price changes are themselves grounded in the newly calibrated belief state. This cycle – real trade $\rightarrow$ gradient update $\rightarrow$ shadow trades – is the core feedback loop of the ParlayMarket.

Relation to Properties 1–3. The three mechanisms above address our desiderata from Section 3 in turn. We now give an informal preview of the reasons behind this, and then prove formal results and empirical justification in the succeeding sections.

Property 1 (Liquidity consistency) is satisfied because all N_M posted prices flow from a single shared φ , there is no separate model or order book per market. More importantly, shadow trades mean the AMM does not need direct trading activity in a market to keep its price current: every real trade on any correlated market triggers updates that ripple through φ and shift all posted prices simultaneously. A rarely-traded parlay market is not forgotten; it benefits from every trade on its correlated counterparts. Complete proof is deferred to appendix of the full version.

Property 2 (Bounded loss) follows from the balance that shadow trades create between gradient signal and gradient noise. Each real trade adds one update to φ , but the shadow trades it spawns spread that update’s effect across all related markets. As the number of markets grows, both the number of updates per round and the noise in each update grow together and cancel, keeping the steady-state parameter error, and therefore the AMM’s total monetary loss per round, bounded independently of M . Without shadow trades this cancellation breaks, and loss grows exponentially with M .

Property 3 (Information aggregation) follows from the structure of the cross-entropy objective. The composite pseudo-likelihood $\mathcal{L}(\varphi)$ has a unique minimum at the parameters φ^* whose implied moments match those of the true distribution P^* . Because each trade gradient pushes φ toward this minimum, the AMM’s posted prices converge to those of the best Ising approximation to P^* as trade flow accumulates, and with it, the posted joint probabilities approach the true joint distribution.

5.2 Convergence to the True Distribution

We first show how Property 3 is satisfied by focusing on how well the Ising model learns the joint distribution via the incoming trades. Our main result concerns the posted prices of the ParlayMarket – they converge geometrically to the best Ising approximation of P^* . After $O(1/(\eta\mu))$ informative trades the KL divergence between the true joint distribution and the AMM’s belief shrinks to the model-misspecification floor, and all N_M posted prices stabilize at their moment-matched values.

► **Theorem 1** (Geometric convergence). *In the pure arbitrageur environment, starting from $\varphi_0 = 0$ with step size $\eta \leq 1/(2L)$:*

$$\mathbb{E}\|\varphi_t - \varphi^*\|^2 \leq (1 - 2\eta\mu)^t \|\varphi^*\|^2 + \frac{\eta\sigma^2}{2\mu},$$

where $\sigma^2 = \mathbb{E}\|\mathbf{g}_t(\varphi^*)\|^2$ is the gradient variance at the optimum. With $\eta = 0.2$ and $\mu \approx \frac{1}{4}$, the contraction factor per round is 0.9; roughly 22 rounds suffice to reduce the initial error by $10\times$.

Three results underpin this guarantee.

Moment-matched target. The fixed point φ^* is the unique parameter vector satisfying $p_i^{\varphi^*} = p_i^*$ for all i and $p_{ij}^{\varphi^*} = p_{ij}^*$ for all $i < j$. This follows directly from the update rule: at φ^* the per-trade scalar $\lambda_m = (p_m^{\varphi^*} - p_m^*)/(1 - p_m^{\varphi^*})$ vanishes for every market m , so all gradient steps are zero. Equivalently, φ^* is the I-projection of P^* onto the Ising family – the maximum-entropy distribution whose first and second moments match those of P^* [25]. Once $\varphi_t \rightarrow \varphi^*$, the posted joint prices $p_S^{\varphi} = P_{\varphi}(\bigcap_{i \in S} \{X_i = 1\})$ converge to the best Ising-family approximation of P^* , giving $D_{\text{KL}}(P^* \| P_{\varphi_t}) \rightarrow 0$ up to the misspecification gap from higher-order dependence not captured by the pairwise family.

Loss curvature. The contraction rate $(1 - 2\eta\mu)$ is determined by the strong-convexity constant μ of the composite pseudo-likelihood $\mathcal{L}$.

► **Proposition 2** (Strong convexity). *On the compact ball $\Phi_B = \{\varphi : \|\varphi\| \leq B\}$, $\mathcal{L}$ is μ -strongly convex. At the uninformative initialisation $\varphi = 0$, the Hessian is diagonal with $\mu = \frac{3}{16}$ (interaction directions) and $\Lambda = \frac{1}{4}$ (bias directions), giving condition number $\kappa = \Lambda/\mu = \frac{4}{3}$. All eigenvalues are $O(1)$, independent of M .*

Each CE term in $\mathcal{L}$ contributes a positive-semidefinite rank-one block

$$H_m = \frac{(\nabla_{\varphi} p_m^{\varphi})(\nabla_{\varphi} p_m^{\varphi})^{\top}}{p_m^*(1 - p_m^*)}$$

to the Hessian via the covariance identity. At $\varphi = 0$ the Ising model is the uniform product measure, so cross-terms vanish and the Fisher matrix is diagonal with entries $\frac{1}{4}$ (bias parameters) and $\frac{3}{16}$ (interaction parameters) – all $O(1)$ regardless of M . Adding more base markets adds more CE terms and therefore more curvature, but the per-parameter curvature stays bounded, which is why the convergence rate does not degrade as M grows.

Parlay acceleration. Higher-order parlay trades strictly improve the convergence rate by adding curvature in the interaction directions that base-market trades cannot supply alone. When the AMM accepts k -way parlay bets for $k \geq 3$, each parlay- S trade contributes a further PSD block H_S to the Hessian, raising the effective strong-convexity constant to

$$\mu_{\text{eff}} \geq \mu + \sum_{k \geq 3} \sum_{S \in \mathcal{S}_k} \lambda_S \mu_S, \quad \mu_S = \lambda_{\min}(H_S) \geq 0, \quad (14)$$

with strict inequality whenever parlay S is informative about φ . The contraction factor improves from $(1 - 2\eta\mu)$ to $(1 - 2\eta\mu_{\text{eff}})$, and the cumulative regret tightens to $G^2/(2\mu_{\text{eff}})(1 + \log T) = O(\log T)$. The gain is largest for the interaction weights W_{kl} : a two-way parlay $\{k, l\}$ contributes only $O(\rho^2)$ to the (W_{kl}, W_{kl}) curvature entry via base-price cross-effects, whereas a higher-order parlay $S \ni k, l$ contributes $O(1)$ directly – a factor of $O(\rho^{-2})$ improvement at small correlations ρ . Proofs are provided in Appendix C of the full version.

5.3 Loss Bounds

The previous subsection established Property 3 by showing that the shared parameter vector φ_t converges to the moment-matched target φ^* . We now turn to Property 2. The question here is not whether the AMM learns the correct joint distribution, but how the remaining parameter error translates into monetary loss. The bridge is a quadratic sensitivity bound: near the fixed point, LMSR loss is second-order in pricing error, and pricing error is first-order in parameter error. In the complete-parlay regime, shadow trades then pool information across all

$$N_M = 2^M - 1$$

active contracts, keeping the *aggregate* steady-state loss per round bounded even though the number of quoted markets grows exponentially.

► **Proposition 3** (Complete-parlay loss scaling). *In the complete-parlay regime, where every non-empty subset $S \subseteq [M]$ is traded at equal rate and the SGD step size is η , the aggregate expected LMSR loss admits the decomposition*

$$\mathbb{E}[L_T] \leq L_{\text{transient}} + T L_{\text{ss,round}},$$

with

$$L_{\text{transient}} = O\left(\frac{b\rho^2 M^2}{\eta}\right), \quad L_{\text{ss,round}} = O(b\eta).$$

In particular, the dependence on M is polynomial rather than exponential, so the complete-parlay Parlay AMM satisfies the bounded-loss requirement of Property 2.

The proposition should be read as a two-phase decomposition. The first term is a one-time learning cost incurred while the AMM moves from the uninformative initialization toward φ^* . The second is the steady-state cost of continuing to learn from a stochastic stream of informative trades once the model has essentially converged.

Loss–parameter link. We first convert the parameter-space error from Theorem 1 into a bound on monetary loss.

► **Lemma 4** (Quadratic sensitivity). *For a market S with true probability p_S^* and AMM parameters φ , the expected LMSR monetary loss from one informed trade satisfies*

$$b \text{KL}(p_S^* \| p_S^\varphi) \leq \frac{b \Lambda_S}{2} \|\varphi - \varphi^*\|^2, \quad \Lambda_S := \sup_{\varphi} \frac{\|\nabla_{\varphi} p_S^\varphi\|^2}{p_S^*(1 - p_S^*)}.$$

The lemma is the composition of two local facts. First, for p near p^* ,

$$\text{KL}(p^* \| p) = \frac{(p - p^*)^2}{2p^*(1 - p^*)} + O(|p - p^*|^3).$$

Second, the price error is linear in the parameter error:

$$p_S^\varphi - p_S^* = \nabla_{\varphi} p_S^* \cdot (\varphi - \varphi^*) + O(\|\varphi - \varphi^*\|^2).$$

Combining the two shows that monetary loss is quadratic in parameter error. This is the key bridge from the convergence theory to Property 2.

Signal–noise balance in the complete-parlay regime. In the formal interaction model, one market is selected uniformly at random each round. Over a block of N_M rounds, each of the $N_M = 2^M - 1$ markets is therefore visited once *in expectation*. The cumulative local information gathered over such a block is summarized by the aggregate sensitivity matrix

$$\Sigma_{\text{cycle}} := \sum_{S \neq \emptyset} F_S(\varphi^*), \quad F_S(\varphi^*) := \frac{(\nabla_{\varphi} p_S^*)(\nabla_{\varphi} p_S^*)^\top}{p_S^*(1 - p_S^*)}.$$

Each term $F_S(\varphi^*)$ measures how informative market S is about the shared parameter vector near the truth: markets whose prices respond strongly to perturbations of φ contribute more to learning. Because all markets update the same shared state φ , both the mean

restoring drift toward φ^* and the stochastic gradient noise accumulate through this same pooled matrix Σ_{cycle} . Put differently, adding more markets increases signal and noise at the same rate rather than letting noise dominate. As a result, the steady-state parameter error remains $O(\eta)$ rather than growing with N_M . Applying Lemma 4 then gives

$$L_{\text{ss,round}} = O(b\eta).$$

Why shadow trades matter. This favorable scaling is not automatic: it relies on the fact that a trade in one market updates the shared parameter vector and therefore refreshes all related quotes. That cross-market propagation is exactly what shadow trades provide. A trade in market S does not only improve the quote for S ; it also carries information about overlapping base and parlay contracts through the common state φ . Without this mechanism, each additional market would still add its own stochastic noise, but the corresponding information would remain largely confined to the market that was directly traded. The system would then behave more like a collection of weakly connected books than a single pooled market maker, and the uniform aggregate bound in Proposition 3 would no longer be expected to hold.

As an immediate corollary, once the transient learning phase has passed,

$$\bar{\ell}_M \approx \frac{L_{\text{ss,round}}}{N_M} = O\left(\frac{b\eta}{2^M}\right),$$

so the *mean loss per listed market* decays exponentially with M . We state this as a corollary rather than part of Proposition 3 because the design objective in Section 3 is bounded *aggregate* loss; the halving pattern in Figure 3 is the complete-parlay consequence of that aggregate bound.

5.4 Information-theoretic optimality

The quadratic loss scaling in Proposition 3 is worst-case optimal. To see why, consider the pairwise family

$$P_s(x) \propto \exp\left(\gamma \sum_{i < j} s_{ij} x_i x_j\right), \quad s_{ij} \in \{-1, +1\}.$$

This family contains one independent hidden correlation sign for every pair of events, hence $\binom{M}{2} = \Theta(M^2)$ independent pairwise facts. Any AMM that does not know these signs in advance must learn them from order flow in order to price the corresponding two-leg parlays correctly.

For a fixed pair (i, j) , the two possible signs $s_{ij} = +1$ and $s_{ij} = -1$ imply different fair prices for the parlay $X_i X_j$. Until the sign is learned, any AMM quoting this parlay with liquidity scale b must be wrong on one of the two cases, and an informed trader can extract loss from that mispricing. For constant-strength correlations this loss is $\Omega(b)$ per unresolved pair; for weak correlations of size γ it is $\Omega(b\gamma^2)$.

Since there are $\Theta(M^2)$ independent pairwise signs, even the best AMM in this class has worst-case information-acquisition loss

$$\inf_A \sup_s \mathbb{E}_{P_s}[L_T] \geq \Omega\left(b \binom{M}{2}\right) = \Omega(bM^2),$$

for constant-strength correlations, and $\Omega(b\gamma^2 M^2)$ in the weak-correlation regime. Thus the quadratic loss rate is not an artifact of the Ising update rule: it is the unavoidable cost of learning dense pairwise correlation structure from trades. *ParlayMarket* matches this rate up to constants while avoiding the exponential loss of treating each parlay as an independent market.

5.5 Parlay Trades as essential for learning

The previous subsection used Theorem 1 only as an input to the loss bound. We now return to the learning dynamics themselves. Higher-order parlays are not merely extra contracts to quote; they also sharpen the learning problem by adding curvature to the shared objective. This strengthens Property 3 directly, and through Lemma 4 it improves the constants in Property 2 as well.

► **Proposition 5** (Parlay-accelerated convergence). *Suppose, in addition to base and pairwise markets, the AMM receives trades on higher-order parlays $S \subseteq [M]$ with $|S| \geq 3$ at rates $\{\lambda_S\}$. Let $\mathcal{L}^+$ denote the augmented composite loss including these parlay terms, and write*

$$\mathcal{S}_k := \{S \subseteq [M] : |S| = k\}.$$

Then:

1. **Improved local curvature.** *In a neighborhood of φ^* , the augmented loss has local strong-convexity constant μ_{eff} satisfying*

$$\mu_{\text{eff}} \geq \mu + \sum_{k \geq 3} \sum_{S \in \mathcal{S}_k} \lambda_S \mu_S, \quad \mu_S \geq 0,$$

with strict inequality whenever at least one traded higher-order parlay is informative about φ .

2. **Faster parameter convergence.** *Under the same SGD dynamics (13) with step size $\eta \leq 1/(2L)$,*

$$\mathbb{E} \|\varphi_t - \varphi^*\|^2 \leq (1 - 2\eta\mu_{\text{eff}})^t \|\varphi^*\|^2 + \frac{\eta\sigma^2}{2\mu_{\text{eff}}},$$

yielding strictly faster geometric contraction than under the base-and-pairwise objective alone.

3. **Improved average error and loss.** *Let*

$$\bar{\varepsilon}_T^2 := \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\varphi_t - \varphi^*\|^2.$$

Then

$$\bar{\varepsilon}_T^2 \leq O\left(\frac{\|\varphi^*\|^2}{\eta\mu_{\text{eff}}T} + \eta\right).$$

Consequently, the average LMSR loss per round is $O(b\bar{\varepsilon}_T^2)$ and the cumulative excess loss over T rounds is $O(bT\bar{\varepsilon}_T^2)$, both with strictly improved constants.

Intuition: curvature as information aggregation. Each trade provides a stochastic gradient whose strength depends on how informative the observed event is about the parameters. Base-market trades mainly constrain first moments, and pairwise trades constrain interactions only one edge at a time. By contrast, a k -way parlay directly reveals a joint probability and therefore aligns with several interaction directions simultaneously. At the level of the Hessian, each informative parlay adds a positive semidefinite block, strengthening the restoring force toward φ^* .

Interaction-level acceleration. The gain is largest for interaction weights W_{ij} . In weak-correlation regimes, these directions are the hardest to learn from singleton flow alone, because the relevant gradients appear only through small cross-effects. Higher-order parlays that include both i and j constrain several interaction directions at once and can therefore dominate the curvature in the weakly identified subspace; heuristically, the improvement can be of order $O(\rho^{-2})$ when $\rho \ll 1$.

From faster learning to lower loss. Proposition 5 should be read as the mechanism behind the favorable loss scaling in Proposition 3. By Lemma 4, monetary loss is quadratic in parameter error, so any increase in curvature reduces both the transient learning cost and the steady-state error floor. Parlays are therefore not only additional products to quote; they are an additional source of statistical signal that sharpens the learning dynamics of the market maker itself.

6 Evaluation

We evaluate ParlayMarket in both controlled and market-based settings. The synthetic experiments test the theoretical predictions of Section 5: bounded aggregate loss, geometric convergence, and the role of parlay order flow in accelerating learning. We then study a historical-market replay on Kalshi data to evaluate whether the same advantages persist under observed prices, RFQ flow, and realistic implementation constraints.

6.1 Complete-Parlay Loss Scaling

Setup. We simulate $M \in \{4, 5, 6, 7, 8, 9\}$ correlated binary assets drawn from a jointly Gaussian scoring model with pairwise correlation $\rho = 0.3$. At each round, one of the $N_M = 2^M - 1$ non-empty parlay markets is selected uniformly at random and an informed trader submits a bet of size b . The AMM updates its Ising parameters via SGD with step sizes $\eta_h = \eta_J = 0.2$. All results are averaged over 100 independent simulation runs.

Baselines. We compare the Correlation AMM against four baselines. The *true oracle* knows P^* exactly and serves as a lower bound on achievable loss. The *independent LMSR* maintains a separate LMSR per market, ignoring correlations entirely. The *pairwise oracle* knows the true pairwise marginals p_{ij}^* analytically and fits an Ising MRF to these moments, but does not learn from trade flow. The *Gaussian oracle* uses the true Gaussian copula parameters to price joint events via multivariate normal CDF; it captures higher-order structure exactly under the Gaussian model but cannot update from trades. All oracle baselines are given their parameters for free and serve to bracket the loss achievable with varying degrees of distributional knowledge.

Per-market loss decay. Figure 3 shows the mean per-market LMSR monetary loss ℓ_M as a function of M . Consistent with Proposition 3, ℓ_M decays exponentially as $C \cdot 2^{-M}$, halving with each additional base asset. The product $\ell_M \cdot 2^M \approx 0.15$ is approximately constant for $M = 4, \dots, 9$ (See Table 3 in Appendix A of the full version), confirming that the total per-round loss $N_M \cdot \ell_M$ is essentially flat – consistent with the steady-state dominance predicted by the two-phase decomposition.


(a) Mean LMSR loss per round per market $\bar{\ell}_M$.(b) CVaR₉₅ loss per round per market.

Figure 3 The ParlayMarket is the only non-oracle model whose loss decays exponentially with the number of base markets M , in the presence of parlays; all other non-oracle models plateau or grow.


(a) Mean LMSR loss per round per market $\bar{\ell}_M$.(b) CVaR₉₅ loss per round per market.

Figure 4 The ParlayMarket fails to perform better than even independent AMMs in the absence of parlay markets.

Price convergence. Figure 6 in Appendix A of the full version shows the mean absolute error (MAE) between the AMM’s posted prices and the true probabilities over time, averaged across all N_M markets and all simulation runs, with $\pm 1\sigma$ shading. The top panel shows the LMSR market-price MAE; the bottom panel shows the Ising marginal belief error for the Ising model. Both converge monotonically; the Ising belief error converges faster because belief propagation directly optimises the Ising objective, while the market-price error includes the additional LMSR rounding effect.

Parlay acceleration. Consistent with Proposition 5, the advantage of the ParlayMarket is entirely contingent on the presence of parlay order flow. Figure 3 shows that with complete parlays the Ising AMM achieves the lowest non-oracle loss, decaying exponentially with M . Figure 4 shows the complementary picture: when informed traders are restricted to base markets only, the model ranking completely reverses.

6.2 Importance of Parlays for learning

Setup. We run an ablation in which informed traders are restricted to the M base markets only – no parlay bets are placed – while all AMMs still maintain and price all $N_M = 2^M - 1$ markets. Figures 3 and 4 show mean loss per round per market vs. M for both conditions.

The model ranking reverses completely. With complete parlays (Figure 3), the ParlayMarket is the *best-performing* non-oracle model: it achieves the lowest per-round loss of all non-oracle AMMs, sitting just above the true oracle and well below every other model. The independent LMSR, pairwise oracle, and Gaussian oracle all incur substantially higher and exponentially growing losses as M increases, reaching ≈ 2.24 , ≈ 0.66 , and ≈ 0.93 per round respectively at $M = 9$.

Without parlays (Figure 4), the ranking *completely reverses*: the ParlayMarket now has the *highest* loss of all non-oracle models, while the independent LMSR, pairwise oracle, and Gaussian oracle cluster together at much lower values (≈ 0.02 – 0.06 per round across all tested M). The y-axis scale tells the story: the no-parlay plot peaks at ≈ 0.18 , while the with-parlay plot requires a scale of 2.3 – a $13\times$ difference – just to accommodate the losses of the non-Ising models.

Why the ranking reverses. The reversal reveals precisely what parlays provide: the *gradient signal needed to calibrate the Ising joint belief model*. Without parlays, the ParlayMarket receives gradient updates only from the M base markets. This is insufficient to pin down the $\binom{M}{2}$ interaction parameters W_{ij} in any reasonable simulation length: the base-market signal is under-determined relative to the full Ising parameter space, so the model misprices joint events and incurs high loss. Simpler models – the pairwise oracle (which knows p_{ij}^* analytically) and the independent LMSR (which only needs to price M uncorrelated events) – do not face this calibration problem and therefore achieve low loss without parlays. With parlays, the situation reverses. Each parlay bet supplies a direct gradient signal on a specific joint event, giving the Ising model the rich information it needs to calibrate $(\theta, \mathbf{W})$ from all N_M markets simultaneously (Proposition 5 and Figure 3). Once calibrated, belief propagation ensures that all N_M posted prices are globally consistent with the learned joint distribution, so adversaries cannot exploit mispriced correlations. The independent LMSR cannot use this information at all (it ignores correlations by design), and the pairwise and Gaussian oracles only capture lower-order structure – leaving them vulnerable to adversaries who trade higher-order parlays against their mispriced joint events.

Implication. Parlays are not merely an efficiency accelerant: they are the *primary information channel* through which the ParlayMarket learns the joint distribution. Without them, the Ising model is the worst non-oracle performer; with them, it is the best. No other model tested achieves this reversal, confirming that the ParlayMarket and parlay markets are jointly necessary – neither is sufficient alone.

6.3 Robustness Under Noisy Informed Traders

Setup. We repeat the complete-parlay simulation with the same parameters but vary the noise-trader fraction $\alpha \in [0, 1]$: with probability α the arriving trader submits a uniformly random price signal, and with probability $1 - \alpha$ they are fully informed. Results at $M = 9$ are extrapolated from direct simulations at $M \in \{3, 4, 5\}$ using the clean $M = 5 \rightarrow 9$ loss scaling.

Results. Figure 7 in Appendix A of the full version shows the per-market LMSR loss as a function of noise-trader fraction α for $M = 9$. The ParlayMarket’s loss grows gracefully with α – it remains the best-performing non-oracle model across all noise levels tested – because the noisy gradient signal inflates the steady-state error while leaving the convergence rate largely unchanged. All models cross the break-even line only at high noise fractions, demonstrating robustness to realistic noise levels.




Figure 5 Key per-round statistics (Loss / Step) for three multinomial market configurations. In all cases the Ising AMM (ours) sits between the True Oracle lower bound and the Independent baseline, with the largest relative gain visible in the tail-risk measures (95% VaR and CVaR).

6.4 Extension to Multinomial Markets

Setup. All experiments above assume binary outcomes per base asset. We now evaluate whether the Ising AMM generalises to *multinomial* markets, where each base asset can resolve to one of $K > 2$ outcomes. We run three configurations that vary the outcome structure and the nature of price jumps (e.g. smooth updates vs. discontinuous jumps between outcome clusters). In each case, we compare the Ising AMM against a *True Oracle* baseline (knows the true joint distribution exactly) and an *Independent* baseline (maintains a separate market maker per outcome, ignoring all correlations). The underlying trader model for these results is a straightforward generalization of the Gaussian scoring model. Details of the multinomial model are provided in Appendix D of the full version.

Results. Figure 5 reports the key per-round statistics (mean, median, 95% VaR, and 95% CVaR of Loss / Step) across the three configurations. In every scenario the Ising AMM sits strictly between the True Oracle lower bound and the Independent baseline: it captures a large fraction of the correlation structure available to the oracle while outperforming the disjoint-market-maker strategy on all four statistics. The gap between Ising AMM and Independent is particularly pronounced in the tail-risk metrics (95% VaR and CVaR), indicating that joint modelling not only reduces average loss but also curtails the worst-case exposure that arises when correlated outcomes are priced independently. These results demonstrate that the Ising parameterisation and the parlay-driven gradient signal are not artefacts of the binary setting. The mechanism extends naturally to multinomial markets and retains its loss advantage over disjoint market makers even as the outcome space grows.

6.5 Empirical Market Performance

To evaluate whether these advantages survive outside the synthetic setting, we replay the mechanism on historical Kalshi data. The empirical study focuses on three implementation details that are central to the practicality of the design.

Two execution settings. We consider *two execution settings*. In a *full-market emulation*, base markets are modeled as LMSRs under the AMM, and the informative signal extracted from each trade is its *target price*: an informed trader is assumed to move the market to the price implied by her information. In a *standalone market-maker emulation*, base markets remain external and the AMM is used only as a shared quoting and underwriting engine for

parlays; in that case the relevant signal is *order flow* – trade direction and share quantity – rather than the final market price. These two settings correspond to two distinct practical roles of the mechanism: a complete market design and a drop-in correlation-aware parlay market maker layered on top of an existing RFQ-style platform.

Real-market emulation. The replay is performed on a *real historical universe*. We use Kalshi’s NBA slate for 2026-03-07, comprising 732 base markets across six games and 5,918 listed combo tickers with legs entirely inside this universe, spanning 2- through 10-leg parlays. To scale inference to clusters of this size, we train one Ising model per game cluster and use *loopy belief propagation* rather than exact enumeration. Each model is initialized from the first 30 candle mid-prices of its constituent legs, with fields initialized from median log-odds and couplings initialized at zero, after which the stream of candle updates and parlay trades is replayed chronologically. Candle updates train local fields, while parlay trades update both fields and interactions through the trade-implied target price. The replay therefore evaluates performance on historically accepted RFQ flow rather than on the full distribution of submitted RFQ requests.

Preprocessing. The empirical setup explicitly addresses a key market imperfection: *fee bias*. Observed Kalshi combo execution prices are ask quotes rather than direct probability assessments, so they are systematically upward biased by platform markup. Training directly on these prices would therefore bias the learned joint model upward. To correct for this, we preprocess each observed trade into a canonical all-YES representation using the complement rule and inclusion–exclusion; for mixed YES/NO parlays, the additive fee cancels in the alternating expansion, yielding an approximately unbiased target in the model’s native parameterization. This correction is essential in practice: without it, parlay loss remains flatter and revenue is lower.

The full replay pipeline and implementation details are provided in Appendix E of the full version. Table 2 shows that standalone ParlayMarket market maker achieves the strongest risk-adjusted liquidity-provider portfolio, with the highest Sharpe ratio among the non-native strategies while still accepting a large fraction of historical flow. The full AMM accepts all trades but exhibits lower ROI, likely because the observed RFQ flow contains substantial noise that a fully internalized LMSR-style implementation must absorb directly. Taken together, these results show that the empirical value of ParlayMarket pricing is strongest in the *standalone market-maker* regime: the model can be inserted into an existing platform, improve LP portfolio quality, and remain competitive on accepted quotes.

6.6 Summary

Across both synthetic and historical evaluations, the results support the same conclusion. A shared, learned representation of the joint distribution allows ParlayMarket to quote a combinatorially large family of contracts while controlling aggregate loss, using parlay order flow as a direct signal for dependence learning. In synthetic settings this produces the predicted scaling and convergence behavior; in historical replay it yields a practically useful correlation-aware quoting engine with strong risk-adjusted performance that can be readily deployed.

■ **Table 2** Emulation on the March 7, 2026 NBA slate under the market-competitive filter.

Strategy	Accepted Trades	MM Win Rate	Net PnL	Sharpe
Kalshi (execution price)	5436	84.5%	\$244,190.77	0.7151
Standalone MM (θ_{corr})	3973	86.8%	\$113,984.58	1.0049
Full AMM	5436	69.0%	\$158,577.66	0.7062

7 Conclusion and limitations

In this paper, we showed that parlay trading can be incorporated into an automated market maker without requiring a separate pricing and risk-management process for every combination. The mechanism supports 2^M base and joint contracts using only an $O(M^2)$ shared representation, and correspondingly achieves $O(M^2)$ aggregate information-retrieval loss rather than loss that scales with the full combinatorial market space. This shared structure also allows parlay order flow to serve as a direct signal for learning dependence structure. The empirical results further suggest that the mechanism is not only theoretically well-founded but close to deployable in practice: historical-market replay on Kalshi data shows that, in the standalone market-maker setting, correlation-aware quoting remains competitive under observed prices, RFQ flow, and realistic execution constraints.

An important scope condition of the current framework is that the traded contracts are assumed to already be expressed in a structurally coherent binary form. In many practical settings, this representation must itself be constructed before correlation-aware pricing becomes meaningful. For example, families of mutually exclusive outcomes may satisfy constraints such as $\sum_{i=1}^k X_i = 1$, even though the underlying latent process remains stochastic. In such cases, a separate preprocessing or structural-learning layer is needed to map raw market data into variables on which the assumptions of the AMM hold. We illustrate one such preprocessing step in the evaluation section. This suggests a broader architecture for automated parlay markets: a structural layer that identifies and normalizes the relevant event representation, followed by a pricing layer such as ParlayMarket that provides automated execution and, as a secondary benefit, elicits dependence from the resulting trade flow.

References

- 1 Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization. In *ACM EC*, 2013.
- 2 Jacob Abernethy, Sindhu Kutty, Sébastien Lahaie, and Rahul Sami. Information aggregation in exponential family markets. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.
- 3 Guillermo Angeris, Tarun Chitra, Alex Evans, and Stephen Boyd. Improved price oracles: Constant function market makers. *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 2020. [arXiv:2003.10001](https://arxiv.org/abs/2003.10001).
- 4 Julian Besag. Statistical analysis of non-lattice data. *The Statistician*, 24(3):179–195, 1975.
- 5 Guy Bresler. Efficiently learning ising models on arbitrary graphs. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, 2015. doi:10.1145/2746539.2746631.
- 6 Yiling Chen, Nicolas Lambert, David M. Pennock, and Jennifer Wortman Vaughan. Eliciting predictions and recommendations for decision making. *ACM Transactions on Embedded Computing Systems*, 9(4):39:1–39:21, 2010. URL: <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/04/TEAC-final1.pdf>.

- 7 Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. *Journal of Artificial Intelligence Research*, 2007.
- 8 Yiling Chen, David M. Pennock, Michael H. Bowling, and Nicolas Lambert. Price updating in combinatorial prediction markets with bayesian networks. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, 2008. [arXiv:1202.3756](#).
- 9 Free Systems. Building the truth machine. <https://open.substack.com/pub/freesystems/p/building-the-truth-machine>, 2024. Substack.
- 10 Rafael Frongillo, Ian Kash, and Mark D. Reid. Interpreting prediction markets: A stochastic approach. In *Advances in Neural Information Processing Systems*, 2012. URL: <https://papers.nips.cc/paper/4529-interpreting-prediction-markets-a-stochastic-approach>.
- 11 Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. Technical report, George Mason University, 2002.
- 12 Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003. doi:10.1023/A:1022058209073.
- 13 Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 2007.
- 14 Elad Hazan. *Introduction to Online Convex Optimization*. Now Publishers, 2019. [arXiv:1909.05207](#).
- 15 Sharif Hossain, Jenn Wortman Vaughan, Yuan Wang, and Yiling Chen. Designing automated market makers for combinatorial securities: A geometric viewpoint. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*, 2012.
- 16 Kalshi. Kalshi: Event contracts exchange. <https://kalshi.com>, 2024. Accessed: 2026-03-19.
- 17 Kathryn Blackmond Laskey, Tai Sing Lee, Robin Hanson, Charles Twardy, Shun-Yang Matsumoto, and Brandon Goldfedder. Graphical model market maker for combinatorial prediction markets. In *Proceedings of the AAAI Workshop on Trading Agent Design and Analysis*, 2012.
- 18 Bruce G. Lindsay. Composite likelihood methods. In N. U. Prabhu, editor, *Statistical Inference from Stochastic Processes*, volume 80 of *Contemporary Mathematics*, pages 221–239. American Mathematical Society, Providence, RI, 1988.
- 19 Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. *arXiv preprint*, 2022. [arXiv:2208.06046](#).
- 20 Abraham Othman and Tuomas Sandholm. When do markets with simple traders fail? In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2010. URL: https://www.ifaamas.org/Proceedings/aamas2010/pdf/01%20Full%20Papers/18_02_FP_0266.pdf.
- 21 Polymarket. Polymarket: A decentralized information market platform. <https://polymarket.com>, 2024. Accessed: 2026-03-19.
- 22 Pradeep Ravikumar, Martin J. Wainwright, and John D. Lafferty. High-dimensional ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010. doi:10.1214/09-AOS691.
- 23 Dan Robinson and Ciamac C. Moallemi. pm-amm. <https://www.paradigm.xyz/2024/11/pm-amm>, 2024.
- 24 Marc Vuffray, Sidhant Misra, Andrey Y. Lokhov, and Michael Chertkov. Interaction screening: Efficient and sample-optimal learning of ising models. In *Advances in Neural Information Processing Systems*, 2016. [arXiv:1605.07252](#).
- 25 Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008. doi:10.1561/22000000001.
- 26 Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, 2003. URL: <https://www.cs.cmu.edu/~maz/publications/techconvex.pdf>.